

Appln. No. 10/763,734
Amendment dated June 12, 2006
Reply to Office Action mailed February 10, 2006

REMARKS

Reconsideration is respectfully requested.

Claims 1 through 19 remain in this application. No claims have been cancelled. No claims have been withdrawn. Claim 20 has been added.

Parts 1 through 4 of the Office Action

Claims 1 through 7, 11 and 13 through 16 have been rejected under 35 U.S.C. §102(e) as being anticipated by Oh.

Claims 8 and 9 and 17 have been rejected under 35 U.S.C. Section 103(a) as being unpatentable over Oh in view of Ebata.

Claims 10 and 18 have been rejected under 35 U.S.C. Section 103(a) as being unpatentable over Oh in view of Wells.

Claims 12 and 19 have been rejected under 35 U.S.C. Section 103(a) as being unpatentable over Oh in view of Watkins.

Claim 1 requires "locating an unused portion of a total disk storage space on disk drives of at least two grid computers of the computing grid" and "presenting a portion of the total disk storage space of each of the at least two grid computers *as a single virtual storage drive* on at least one computer" (emphasis added).

In the rejection of claim 1 in the Office Action, it is stated that:

As per claim 1, Oh teaches a method of creating a virtual disk storage (i.e. the extra layer of storage in the memory hierarchy between the main memory and the hard disk) construct using disk storage consolidated from at least two grid computers (i.e. from multiple exporters) of a computing grid, comprising: locating an unused portion of a total disk storage space on disk drives (i.e. partitions) of at least two grid computers of the computing grid; and presenting a portion of the total disk storage space of each of the at least two grid computers as a single virtual storage drive (i.e. the

Appln. No. 10/763,734
Amendment dated June 12, 2006
Reply to Office Action mailed February 10, 2006

aggregated remote memory; the exported memory partition) on at least one computer (i.e. the importer) (e.g. see paragraph [0025]).

Looking to the referenced portion of the Oh published patent application, paragraph [0025]] states (emphasis added):

[0025] FIG. 2 shows the MPD concept for Case 2 above, i.e., the concept of MPD in a paged-memory system by adding an extra layer in the memory hierarchy between the main memory and the hard disk. A machine 10 borrows two memory partitions from a machine 20 over the network. Although only one memory exporter is shown in the figure, the mechanism allows importing multiple partitions from multiple exporters so that the aggregated remote memory can be huge. The exporter's exported memory partition becomes a new layer in the importer's memory hierarchy.

Considering the referenced portion of the Oh patent application, in the context of the remainder of the application, it is submitted that one of ordinary skill in the art would not be led to "presenting a portion of the total disk storage space of each of the at least two grid computers as a *single virtual storage drive* on at least one computer". More specifically, the Oh patent application discusses an "extra layer in the memory hierarchy between the main memory and the hard disk". It is submitted that, rather than suggesting the presentation of a "single virtual storage drive" as required by claim 1, one of ordinary skill in the art would understand the Oh patent application as discussing an invisible "layer" of memory on a machine receiving the memory partition. Further, the Oh patent application discusses this in the context of memory partitions from a single computer, it states that partitions from more than one machine may be employed, but does not disclose how memory partitions from different computers would be presented. The discussion in paragraph [0025] suggests that each "exported memory partition" becomes "a new layer" on the memory importer, suggesting that there is no unified presentation of the memory partitions from different computers, but as separate, distinct layers. It is submitted that this discussion is more likely to lead one of ordinary skill in the art to understand that the "imported memory partitions" remain distinct on the

Appln. No. 10/763,734
Amendment dated June 12, 2006
Reply to Office Action mailed February 10, 2006

importing computer, rather than leading the skilled in the art to "*presenting* a portion of the total disk storage space of each of the at least two grid computers *as a single virtual storage drive* on at least one computer".

It is therefore submitted that the Oh published patent application would not lead one of ordinary skill in the art to the applicant's claimed invention as defined in claim 1, especially with the requirements set forth above, and therefore it is submitted that claim 1 is allowable over the prior art. Further, claims 2 through 12, which depend from claim 1, also include the requirements discussed above and therefore are also submitted to be in condition for allowance. Similarly, claim 13 requires "*presenting a portion of the total disk storage space of each of the at least two grid computers as a single virtual storage drive on at least one computer*", and for reasons similar to those discussed above, it is submitted that claim 13, as well as claims 14 through 19 which depend from claim 13, are also in condition for allowance.

Claim 3 requires "*reserving a portion of the total disk storage space on each of the at least two grid computers for local use*". Claim 15 similarly requires "*implementing the step of reserving a portion of the total disk storage space on each of the at least two grid computers for local use*". It is alleged in the rejection of the Office Action that:

... Oh teaches that only the unused memory partitions are assigned/contributed to the other/remote computers/grids (e.g. see paragraph [0020]). Therefore, at least a portion (i.e. the used memory partition) of the total disk storage space are inherently reserved for local use on each of the at least two grid computers.

Looking to paragraph [0020] of Oh, it states:

[0020] Memory Partitions on Demand (MPD) is a remote memory allocation mechanism that allows computers in an infrastructure to contribute their unused memory partitions to other computers. The idea is similar to contributing CPU cycles of under-utilized computers in load-balancing of GRID architecture and Internet-based distributed computing, but in this case the shared resource is the main memory.

Appln. No. 10/763,734
Amendment dated June 12, 2006
Reply to Office Action mailed February 10, 2006

However, nothing here states or suggests that there is any reservation of any space on the computers, and to the contrary there is only discussion of "contribut[ing] their unused memory partitions to other computers". It is submitted that this discussion is more likely to lead one of ordinary skill in the art to understanding that there is no restriction or limitation on the "unused memory partitions" of the computer, and that as long as the memory is "unused", the memory is available to be shared. This is in conflict with the requirement of claim 3, which requires that a portion of the total disk storage space is reserved. The rejection says "at least a portion (i.e. the used memory partition) of the total disk storage space are inherently reserved for local use on each of the at least two grid computers". However, one of ordinary skill in the art understands that simply because the computer uses some of its own memory does not mean that there is any memory reserved for its own use. Again, paragraph [00200] refers to "unused" memory, which suggests that any memory that is not used is available to be used by the Oh system. The present invention, and this claim requirement, recognizes that some of the "unused" memory of the grid computer (in addition to the memory in use by the grid computer) may need to be reserved for future usage by the grid computer. The Oh system does not recognize such a reservation.

Claim 4 requires "determining the total disk storage space on each of the at least two grid computers and allocating the total disk storage space between a portion made available for use as part of the virtual storage drive and a portion reserved for local use on the grid computer" (all emphasis added), and the Office Action relies upon the same paragraph [0020]. As noted above, the Oh patent simply does not disclose this feature of the invention.

Claim 5 requires "maintaining a table of grid computers contributing storage space to the virtual storage drive and corresponding amounts of storage space made available by each contributing grid computers". Claim

Appln. No. 10/763,734
Amendment dated June 12, 2006
Reply to Office Action mailed February 10, 2006

16 requires "implementing the step of monitoring at least one of the grid computers for activity indicating that additional disk storage space has been added to the at least one grid computer". The Office Action states that:

... Oh teaches that when a remote computer/grid needs more memory, it can request for it over the network and the remote computer(s)/grid(s) allows to use their unused memory partitions. These remote memory partitions are returned to the original owner when no longer needed by the requesting computer/grid (e.g. see paragraphs [0020]-[0022]). Since the remote memory partitions are returned to the owner grids when they no longer needed, there has to be a table, of grid computers contributing storage space to the virtual storage drive and corresponding amounts of storage space made available by each contributing grid computers, present in the method taught by Oh. Therefore, the feature of having this table is inherently embedded in the method taught by Oh.

The paragraph [0020] has been discussed above, and does not mention any table of "a table of grid computers contributing storage space to the virtual storage drive and corresponding amounts of storage space made available by each contributing grid computers". Further, paragraphs [0021] and [0022] state (emphasis added):

[0021] When a computer in the infrastructure needs more memory, it can request memory partitions over the network and "attach" the remote memory partitions to the local memory. Once attached, the newly added partitions can function in two ways, depending on the situation:

[0022] Case 1. Attaching imported memory directly as an extension of local memory. The attached remote memory partitions function as if more physical memory has been added to the local system (FIG. 1). The additional memory partitions are returned to the original owner when no longer necessary. Page tables, other necessary addressing schemes, and related data structures are properly updated so that a software daemon process "fools" the CPU as if more memory has been added. It is, of course, significantly slower to access the remote memory as compared to the local memory. This delay is justified, as is addressed later below.

This portion of Oh, like the preceding portion, also does not mention any use of a table of grid computers and a table of corresponding amounts of storage space. Although the remarks in the rejection argue that the Oh system must inherently include a table, it is submitted that that is not

Appln. No. 10/763,734
Amendment dated June 12, 2006
Reply to Office Action mailed February 10, 2006

necessarily the case. For example, one of ordinary skill in the art could understand that the Oh system merely collect identifications of computers that have available storage, but not keep track of the amount of available storage, and simply leave it up to peer to peer communications between the memory-requesting and the memory-providing computers to establish the amount of memory available. This, it is submitted that it is not inevitable, or inherent, that the Oh system includes a table meeting the two requirements of claim 5.

Claim 6 requires "monitoring at least one of the grid computers for activity indicating that additional disk storage space has been added to the at least one grid computer", and further claim 7 requires "allocating disk storage space on the at least one grid computer after detecting activity indicating that additional storage space has been added to the at least one grid computer". (Claim 16 requires "implementing the step of monitoring at least one of the grid computers for activity indicating that additional disk storage space has been added to the at least one grid computer".) The benefit to these aspects of the invention is that the system is able to detect and utilize additional memory that is added to the grid computers. The rejection of the Office Action states that:

... Oh teaches that the first and second computers are connected via a network; creating a file in the storage memory of the first computer and allocating additional blocks of the storage memory of the second computer (e.g. see paragraph [0008]). Since Oh teaches that the allocation of additional blocks from the second computer to the first computer based on needs of the first computer, the steps of (i) monitoring at least one of the grid computers (i.e. the first computer) for activity indicating that additional disk storage space (i.e. the additional blocks) has been added to the at least one grid computer, and (ii) allocating disk storage space on the at least one grid computer (i.e. the first computer) after detecting activity indicating that additional storage space has been added to the at least one grid computer, are inherently present in the method taught by Oh.

It is submitted that nothing in the Oh patent application publication that would lead one of ordinary skill in the art to understand that there is any

Appln. No. 10/763,734
Amendment dated June 12, 2006
Reply to Office Action mailed February 10, 2006

monitoring by the Oh system of additional disk storage space being added to one of the grid computers. The detection of activity that indicates that additional storage space has been added to a grid computer is different than simply allocating additional memory to a computer, or detecting that a computer has additional memory needs. Further, paragraph [0008] of Oh states (emphasis added):

[0008] According to an embodiment of the invention, a method for implementing a single address space view for files stored in various distributed systems includes providing first and second machines with first and second storage memories, respectively, wherein the first and second machines are connected via a network; creating a file in the storage memory of the first machine and allocating additional file blocks of the file to the storage memory of the second machine; and establishing a bi-directional file header in the first and second storage memories that links to the file and the additional file blocks.

Again, the discussion here suggests to one of ordinary skill in the art that existing memory is being allocated among connected machines, and not “additional disk storage space being *added* to” one of the computers. It is submitted that one of ordinary skill in the art recognizes the difference between adding additional disk storage space to a computer, and simply allocating memory between networked computers.

Claim 8 requires “monitoring at least one of the grid computers for activity indicating that a minimum amount of free disk storage space has been violated”, and claim 9 requires “allocating disk storage space on the at least one grid computer after detecting activity indicating that the minimum amount of free disk storage space has been violated to restore at least the minimum amount of free disk storage space”. Similarly, claim 17 requires “implementing the step of monitoring at least one of the grid computers for activity indicating that a minimum amount of free disk storage space has been violated”. It is conceded in the Office Action that:

However, Oh failed to teach the further limitations of (i) monitoring at least one of the grid computers for activity indicating that a minimum amount of free disk storage space has been violated; and (ii) allocating disk storage space on the at least one grid computer after

Appln. No. 10/763,734
Amendment dated June 12, 2006
Reply to Office Action mailed February 10, 2006

detecting activity indicating that the minimum amount of free disk storage space has been violated to restore at least the minimum amount of free disk storage space.

It is then contended that:

Ebata, on the other hand, teaches a method for moving files between storages across the network to rebalance the free disk space across the network. Ebata teaches the method includes the step of monitoring at least one of the grid computers (i.e. at least one of the storage across the network) for activity indicating that a minimum amount of free disk storage space has been violated; and (ii) allocating disk storage space on the at least one grid computer after detecting activity indicating that the minimum amount of free disk storage space has been violated to restore at least the minimum amount of free disk storage space (e.g. see the abstract). Accordingly, it would have been obvious to one ordinary skilled in the art at the time of the current invention was made to implement the steps taught by Ebata in the method taught by Oh. In doing so, (i) a steady imbalance of the free disk spaces among the network storages is prevented so that clients can always use the system and even if client writes large files and a maximum quantity of data can be written to disks managed by the virtualized network storage system; and (ii) during file migration between network storages, access requests from clients are not stopped while a file is being moved between network storages (e.g. see paragraphs [0015]-[0016]).

Looking to the Abstract of the Ebata published patent application, it is stated that (emphasis added):

A method for moving files between network storages in a virtualized network storage system having multiple network storages and thereby rebalancing the free disk spaces therein. It periodically judges whether to start rebalancing. Two different conditions are adopted for the determination of start: (a) there is an imbalance of free disk spaces and the minimum value of the free disk spaces is less than a threshold and (b) there is an imbalance of free disk spaces and the frequency of accesses to the system is less than a threshold. Rebalancing is carried out until the difference in free disk spaces falls below the threshold to the extent that the maximum execution time of rebalancing will not be exceeded. If a request to write into a file which is being moved from a client takes place during a file moving step, the movement of the file is aborted and the file is deleted from the destination thereof. Then, another file is selected and the file moving step is carried out again. Or, the write request from the client is abandoned and the movement of the file is continued.

Appln. No. 10/763,734
Amendment dated June 12, 2006
Reply to Office Action mailed February 10, 2006

The Abstract of the Ebata patent application thus leads one of ordinary skill in the art to understand that when an imbalance of free disk occurs, and either the minimum free space on a disk is not present or a minimum number of accesses does not occur, that rebalancing should occur. It is submitted that the Ebata Abstract suggests that the primary consideration for rebalancing is an imbalance in free disk space in the Ebata system, and another condition that may be less than a minimum storage space or less than a minimum disk access. Thus, the minimum free space threshold is not really a "violation" of the minimum amount of free space, as it is only considered after an "imbalance" of free space on the systems is detected.

Further, in the referenced paragraphs [0015] and [0016] of the Ebata patent application, it is stated that:

[0015] A first object of the present invention is to provide a method for rebalancing the free disk spaces in a network storage system virtualized into a single file system view with the disk spaces in the network storages thereof uneven, wherein a steady imbalance of the free disk spaces among the network storages is prevented so that clients can always use the system and even if client writes large files, a maximum quantity of data can be written to disks managed by the virtualized network storage system.

[0016] A second object of the present invention is to provide a method for rebalancing the free disk spaces in a network storage virtualized into a single file system view involving file migration between network storages thereof, wherein access requests from clients are not stopped while a file is being moved between network storages.

Again, it is submitted that one of ordinary skill in the art understands that Ebata is suggesting that it is an imbalance of free disk spaces, and not a violation of a minimum free disk space on a system, that is being detected and acted upon here.

Claim 10 requires "providing a safe area on disk storage space of the virtual storage drive, the safe area being kept free of data", and similarly claim 18 requires "providing a safe area on disk storage space of the virtual

Appln. No. 10/763,734
Amendment dated June 12, 2006
Reply to Office Action mailed February 10, 2006

storage drive, the safe area being kept free of data". It is stated in the rejection of the Office Action that:

Oh... failed to teach that the method additionally including providing a safe area on disk storage space of the virtual storage drive, the safe area being kept free of data. Wells, however, teaches about keeping a portion of the memory space free of data to allow the cleanup operation (e.g. see Col. 5, lines 1-6). Accordingly, it would have been obvious to one ordinary skilled in the art at the time of the current invention was made to implement the step taught by Wells in Oh's method. In doing so, data can be temporarily stored at this safe area when (i) data needs to be transferred within the memory space; and (ii) the cleanup operation is required to run.

However, it is submitted that the Wells patent discusses a disk clean up operation that is unrelated to the Oh system of sharing memory, and it is submitted that one of ordinary skill in the art, considering the Oh, would not look to the Wells patent for modification of the Oh system as Wells discusses "keeping a portion of the memory space free" for the purpose of "allow[ing] the cleanup operation", and this does not pertain to the Oh system, which presumably permits the use of memory of another computer for tasks such as the transfer of data "within the memory space" or for the "cleanup operation".

It is therefore submitted that the cited patents, and especially the allegedly obvious combination of Oh and Ebata set forth in the rejection of the Office Action, would not lead one skilled in the art to the applicant's invention as required by claims 10 and 18.

Withdrawal of the §102(e) and §103(a) rejections of 1 through 19 is therefore respectfully requested.

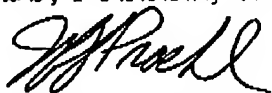
Appln. No. 10/763,734
Amendment dated June 12, 2006
Reply to Office Action mailed February 10, 2006

CONCLUSION

In light of the foregoing amendments and remarks, early reconsideration and allowance of this application are most courteously solicited.

Respectfully submitted,

WOODS, FULLER, SHULTZ & SMITH P.C.



Date: JUNE 12, 2006

Jeffrey A. Proehl (Reg. No. 35,987)
Customer No. **40,158**
P.O. Box 5027
Sioux Falls, SD 57117-5027
(605)336-3890 FAX (605)339-3357

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- ☐ BLACK BORDERS
- ☐ IMAGE CUT OFF AT TOP, BOTTOM OR SIDES
- ☐ FADED TEXT OR DRAWING
- ☐ BLURRED OR ILLEGIBLE TEXT OR DRAWING
- ☐ SKEWED/SLANTED IMAGES
- ☐ COLOR OR BLACK AND WHITE PHOTOGRAPHS
- ☐ GRAY SCALE DOCUMENTS
- ☒ LINES OR MARKS ON ORIGINAL DOCUMENT
- ☐ REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY
- ☐ OTHER: _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.